

# Low False-Alarm Fault Monitoring with Markov Models

Padhraic Smyth, Jeff Mellstrom  
Jet Propulsion Laboratory, MS 238-420  
California Institute of Technology  
4800 Oak Grove Drive, Pasadena, CA 91109.

In this paper we describe a novel approach (based on hidden Markov models) to online fault monitoring which can effectively reduce the false alarm rate of conventional monitoring systems by orders of magnitudes. Continuous monitoring of complex dynamic systems for the purposes of anomaly or fault detection is an increasingly important issue in diverse areas such as onboard vehicle health monitoring, nuclear and chemical plant safety, manufacturing process control, communications network anomaly detection and biomedical health monitoring. In the area of online monitoring of large (34m and 70m) Deep Space Network (DSN) ground antennas (used for tracking deep space spacecraft) we have developed and implemented a unique approach to the problem of detecting and identifying electro-mechanical fault conditions. The underlying theory of the new method is quite general and can in principle be applied to a wide variety of practical problems involving online fault detection in dynamic systems [1]. The key contribution of the Markov approach is the virtual elimination of spurious false alarms which can plague conventional detection systems; this is achieved by taking advantage of the fact that fault conditions typically persist over time. Related applications can be found in [2, 3, 4].

Standard methods for fault detection in dynamic systems have largely relied in the past on two main approaches: control-theoretic methods and heuristic knowledge-based methods. The control-theoretic approach presumes the existence of an accurate and reliable linear system model; frequently in practical applications no such model exists. Knowledge-based methods presume the existence of a domain expert, which may not be true either; furthermore knowledge-based methods are typically not suited for dealing with time-varying systems with feedback.

An alternative approach (which has been proposed in recent years) is based on pattern recognition ideas; a classification model is trained on system data to discriminate between normal and abnormal conditions. Since the model adapts to the data from the system, very little prior knowledge is required, making the method very attractive for many practical situations. However, the characteristics of the system which are used as inputs to the classifier (known as the "features" in pattern recognition jargon) can typically be quite noisy (such as features derived from a typical motor current signal in an electro-mechanical system). This can translate into fluctuations in the decisions of the classifier, resulting in an unacceptably large false alarm rate over time.

In our work at JPL, we have developed a novel system which "smooths" such classification estimates over time using the formalism of hidden Markov models. Hidden Markov models have been used with significant success in speech recognition applications for some time; however it is only with our recent work that their general applicability to fault monitoring has been recognized. Briefly, the system is assumed to follow a Markov trajectory in terms of transiting between normal and abnormal states. The *hidden* aspect of the problem arises due to the fact that these states are not directly observable: only the symptoms (such as sensor time series data) can be directly observed. Hence, the problem is to infer the most likely sequence of system states given the symptom data. In particular, the sampled sensor data is divided into windows over time and from each window a vector of signal parameters is extracted. This data preprocessing step introduces some invariance and robustness with respect to noise. The estimated signal parameter vector provides the input for a standard classification model (we use both kernel density estimators and feedforward neural networks) which in turn produces an instantaneous estimate of the probability of the system's state (conditioned on the particular input). The role of the Markov model is to integrate these probability estimates over time by combining the present state estimate with past state estimates. The hidden Markov model effect is to

dramatically reduce the false alarm rate by several orders of magnitude compared to not using any integration over time. The parameters of the hidden Markov model can be shown to be directly related to the overall mean time between failure of the system and hence can be estimated from prior knowledge of system failure rates.

Pattern recognition systems are typically not sensitive to novel conditions, i.e., they can only detect data from fault classes 011 which the model was trained, data from novel fault classes will be misclassified into one of the known classes. In fault detection this can be a severe problem since the system is quite likely to confront fault conditions which have not been seen before. We have developed an extension to the basic hidden Markov theory which allows the detection of such novel data [5]; furthermore, such data is passed to a fault isolation module which uses a simple linear model to try to infer the exact location of the problem.

The algorithms described above have been successfully tested for online fault detection at a 34m research antenna site. The algorithms were implemented in software via the LABVIEW data acquisition and display package 011 a Macintosh PC. The system produces health estimates every 4 seconds from various sensor signals such as motor current, tachometer voltages, and antenna position readings. In real-time blind tests, involving the introduction of 3 different hardware failures in a controlled manner, the system identified all faults correctly and produced zero false alarms in over 1 hour's worth of monitoring. In comparison a more conventional detection system using Gaussian models, misclassified antenna health 16% of the time and generated 3 distinct false alarms when no faults were present. The Markov monitoring system is currently being integrated for operational use with a new operational 34m DSN antenna at the Goldstone antenna complex. In addition we are currently considering various other applications of the general methodology in applications SUCH as onboard electric vehicle fault monitoring, fault detection in electric rail cars, and online condition monitoring of steel plant machinery.

## Acknowledgements

The research described in this paper was performed at the Jet Propulsion Laboratory, California Institute of Technology, under a contract, with the National Aeronautics and Space Administration and was supported in part by ARPA under grant number N00014-92-J-1860.

## References

1. P. Smyth, 'Hidden Markov models for fault detection in dynamic systems,' *Pattern Recognition*, vol.27, no.1, pp.149-164, 1994.
2. L. P. Heck and J. H. McClellan, 'Mechanical system monitoring using hidden Markov models,' in *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing*, IEEE Press, New York, pp.1697-1700, 1991.
3. E. Ayanoglu, 'Robust and fast failure detection and prediction for fault tolerant communication networks,' *Electronics Letters*, 26(10), pp.940-941, 1992.
4. D. A. Coast, R. M. Stern, G. G. Cano, S. A. Briller, 'An approach to cardiac arrhythmia analysis using hidden Markov models,' *IEEE Trans. Biomedical Engineering*, vol.37, no.9, pp.826-836, 1990.
5. P. Smyth, 'Probabilistic anomaly detection in dynamic systems,' in *Advances in Neural Information Processing Systems 6*, J. D. Cowan, G. Tesauro, J. Alsppector (eds.), Morgan Kaufmann Publishers, 1994, to appear.